Message Latency in Waku Relay with Rate Limiting Nullifiers

Presenter:
Alvaro / Research Engineer at Status.im

Authors:
Alvaro Revuelta
Sergei Tikhomirov
Aaryamann Challani
Hanno Cornelius
Simon Pierre Vivier

Agenda

Background:
- Gossipsub
- Zero-Knowledge and RLN
- Waku

Latency

Simulations
- Single-host simulations
- Multi-host simulations

Contributions

Future work

Questions

Background

Gossipsub:

- Publisher/subscriber based
- Allows a message to reach all peers subscribed to a topic
- In permissionless and adversarial environments
- With constant amplification factor (D)
- Two types of connections: full and metadata
- Eager push + lazy pull approach
- Used in multiple blockchains such as Ethereum
- By Protocol Labs

Background

–

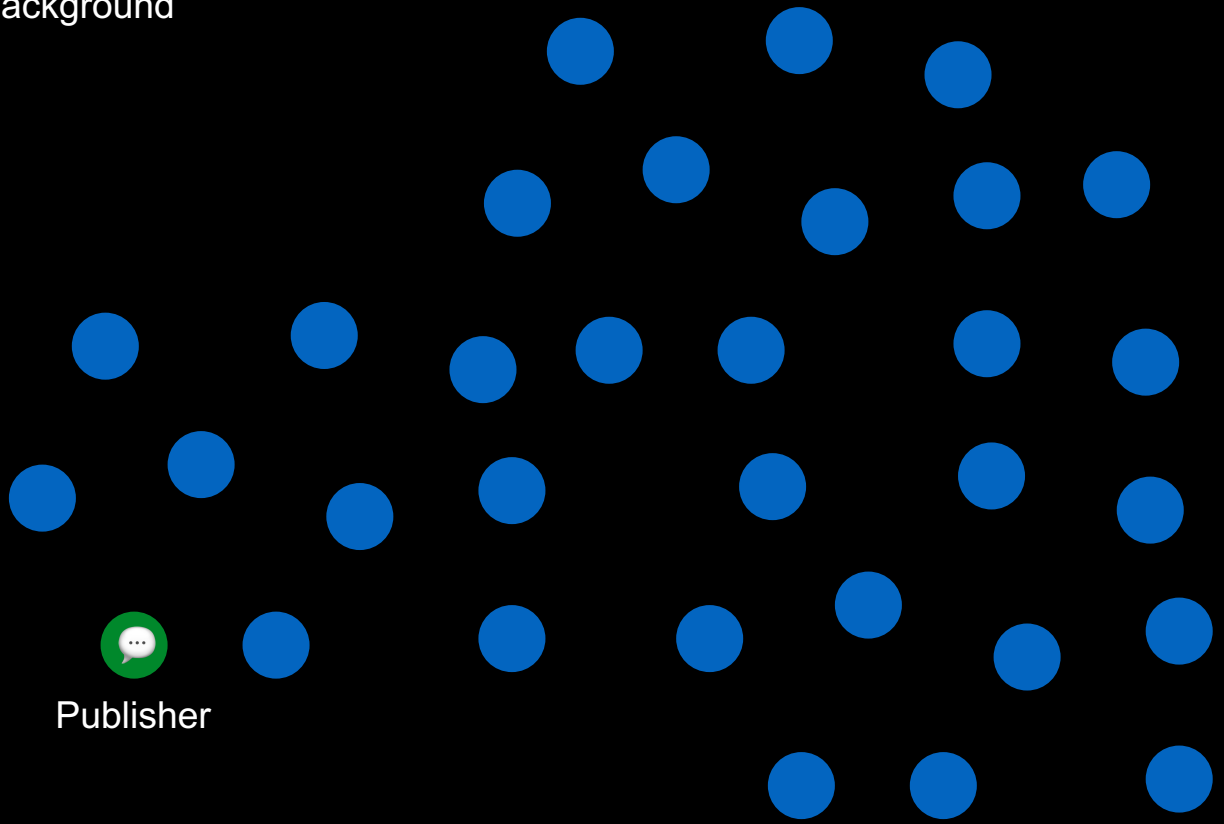

Background

–

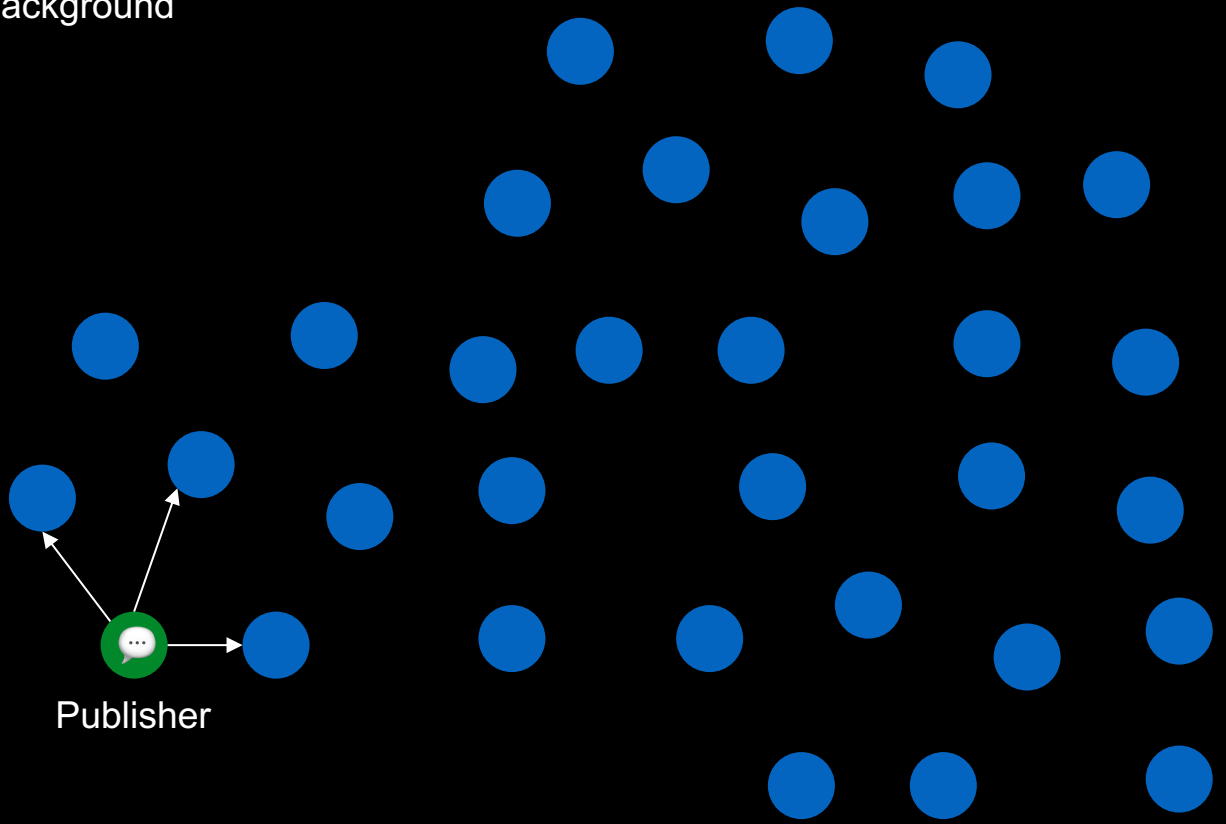

Network D=3
Nodes=32

Background

–

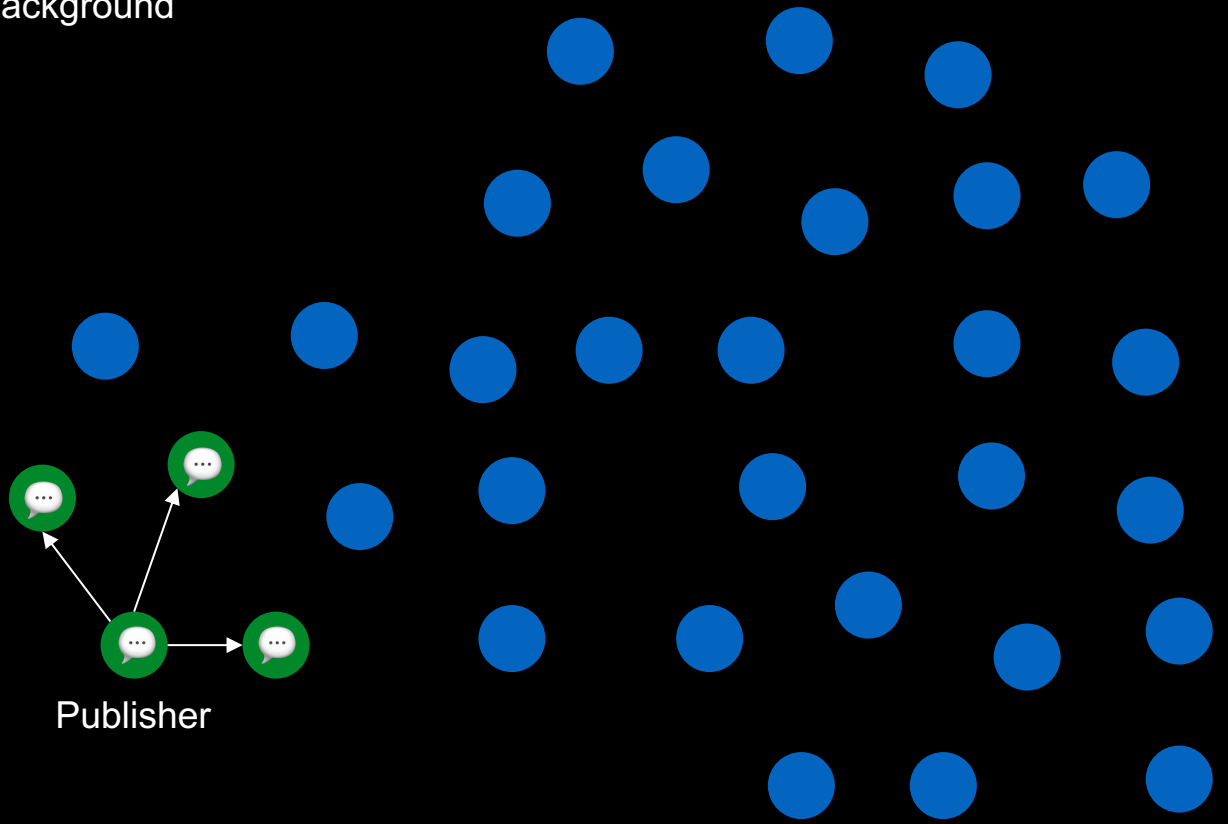

Background

–

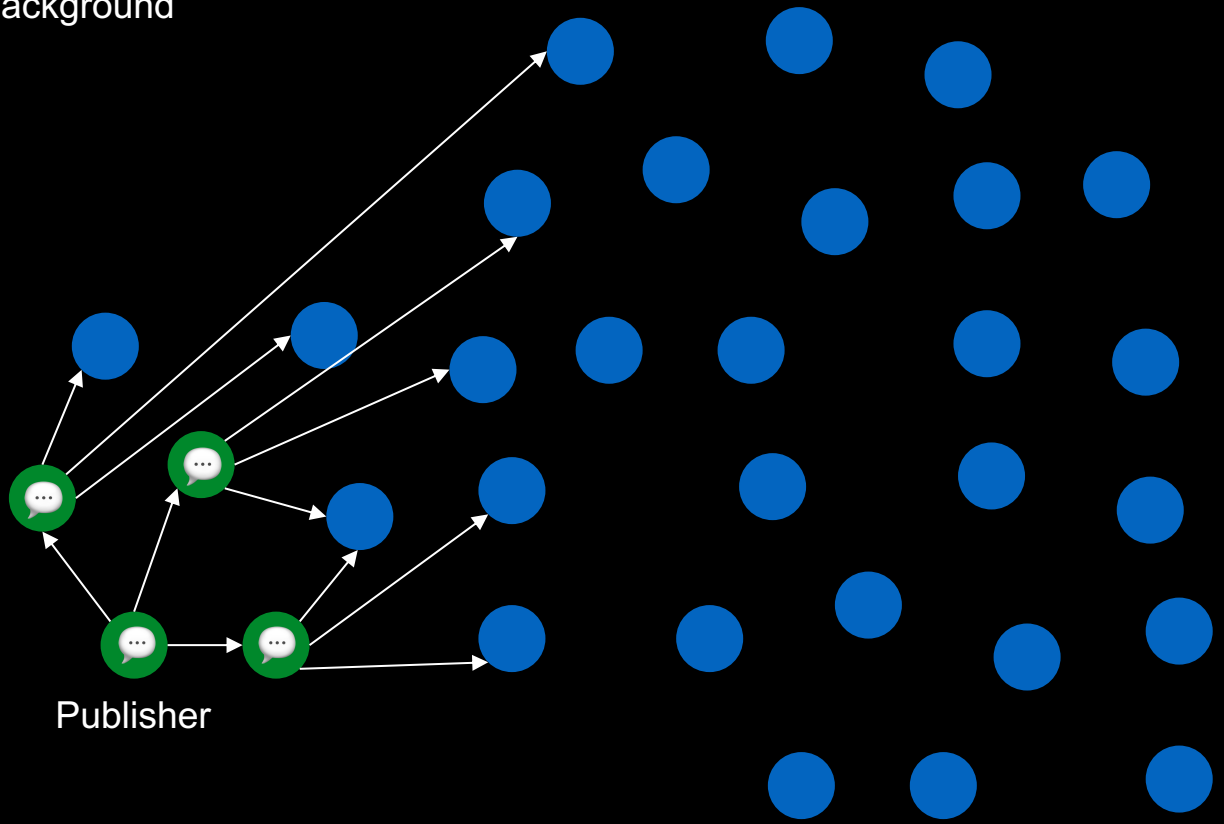

Network D=3
Nodes=32

Background

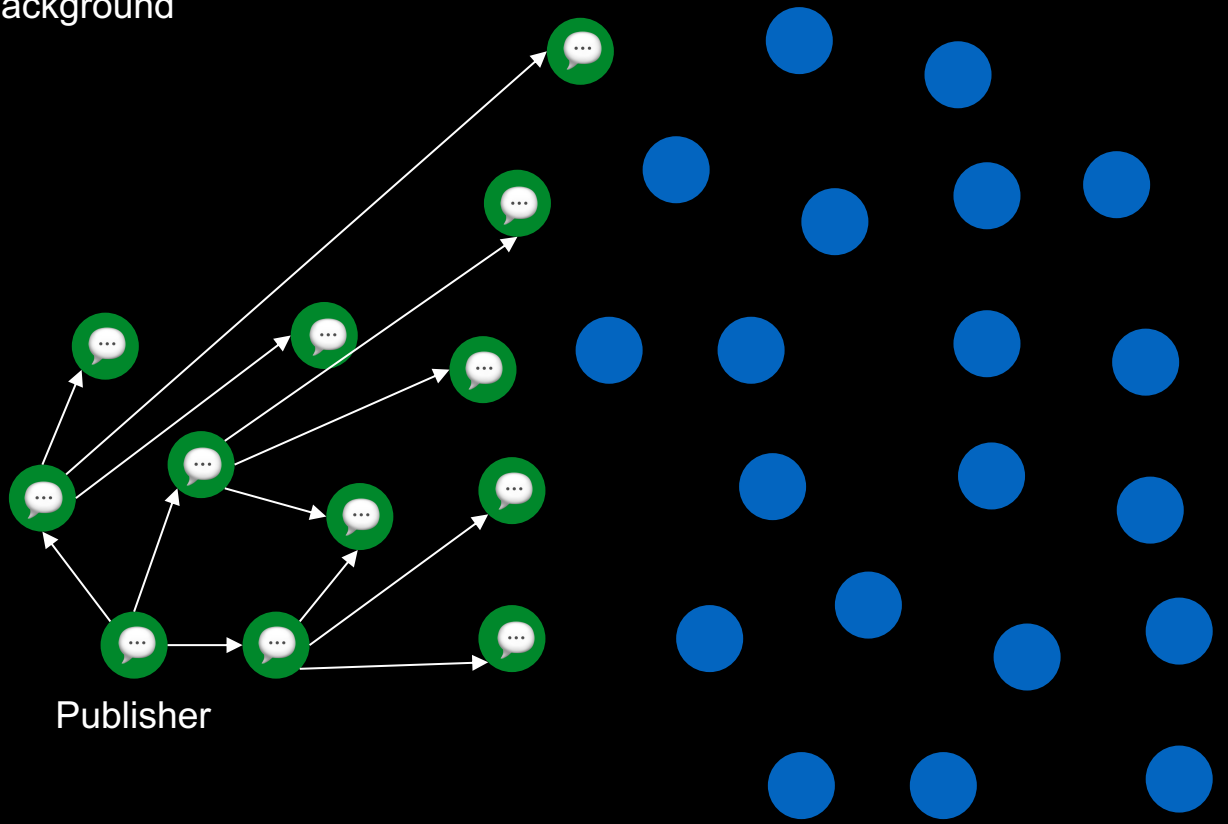

Network D=3
Nodes=32

Background

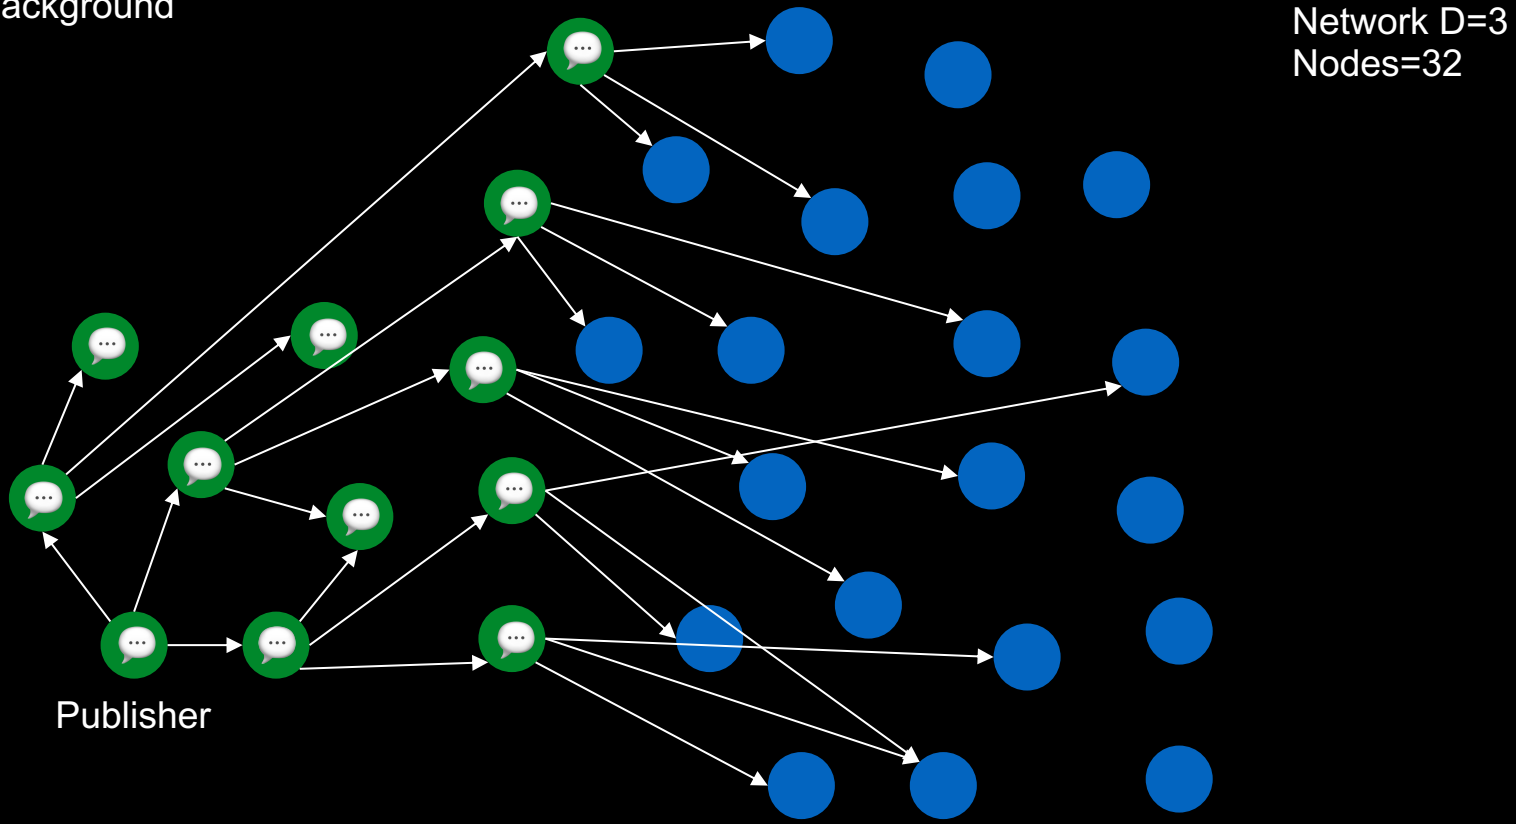

Background

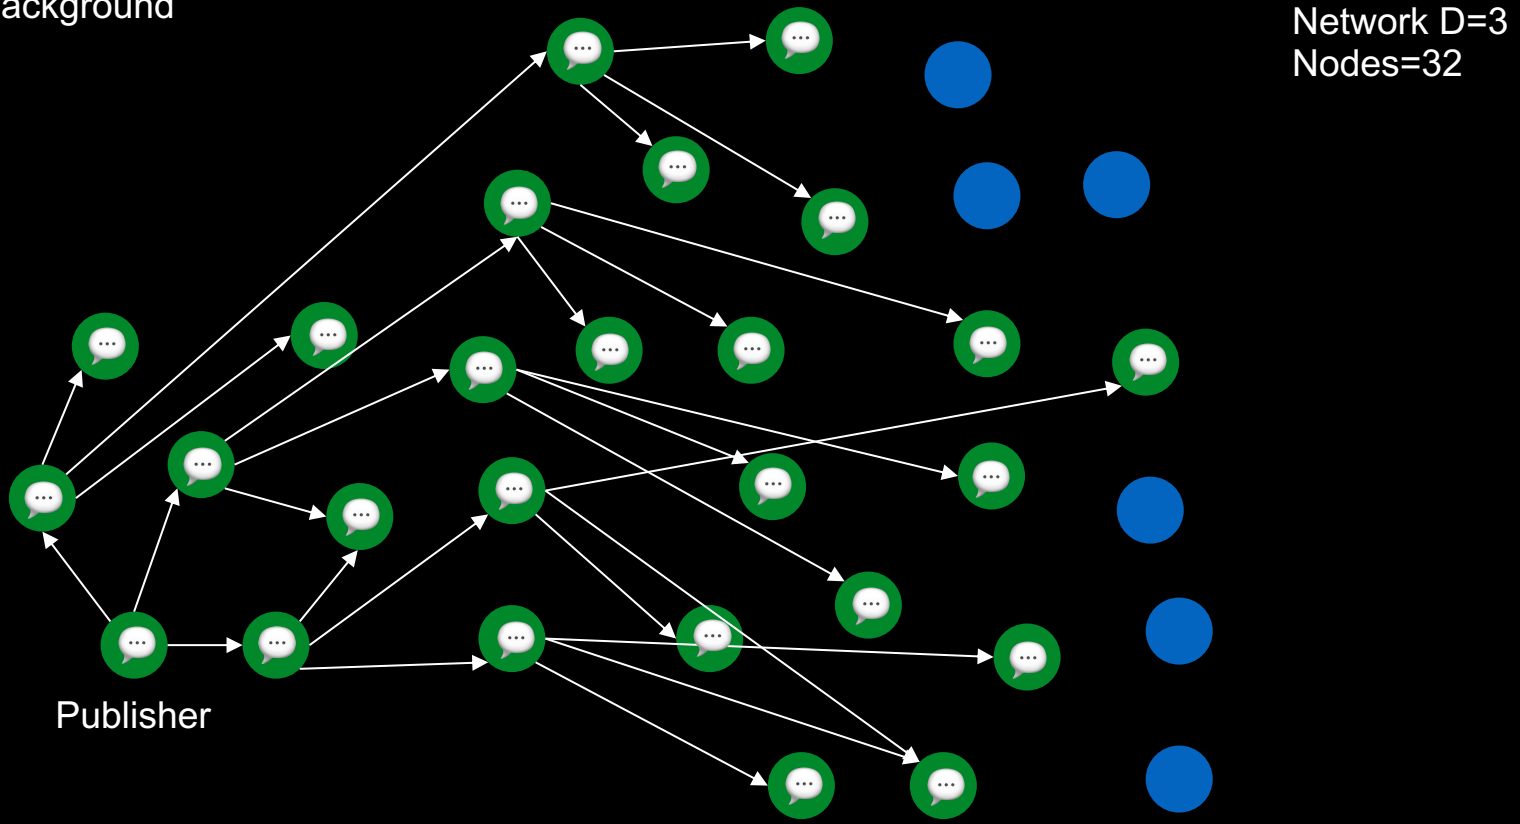

Network D=3
Nodes=32

Background

Network D=3
Nodes=32

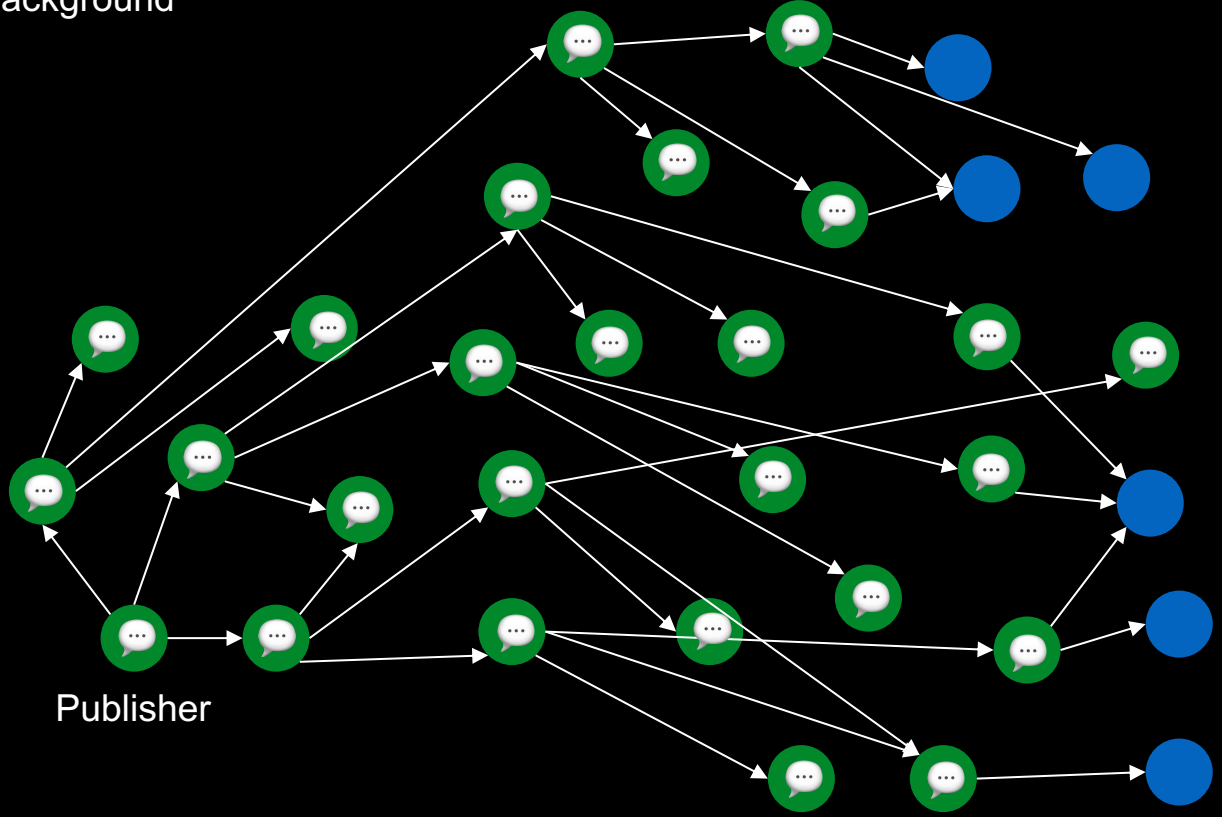

Background

Network D=3
Nodes=32

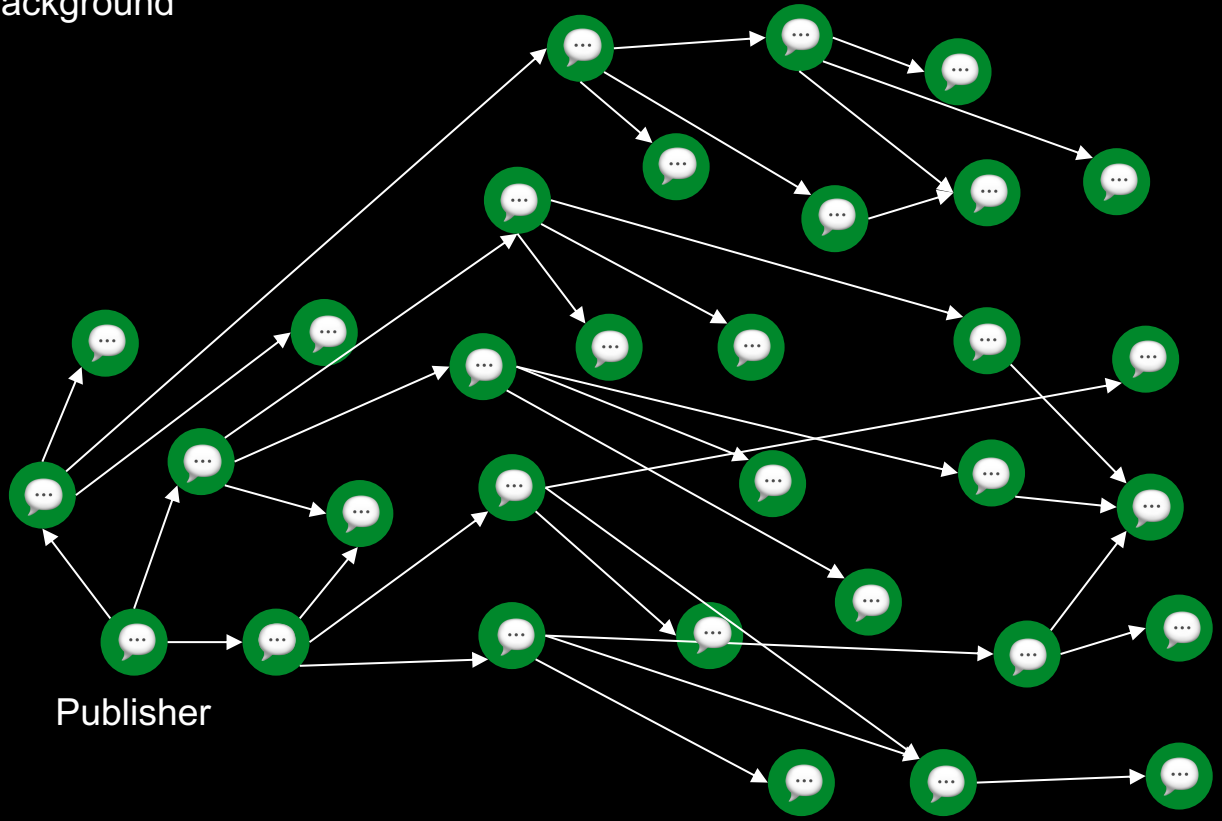

Background

Gossipsub:

- In every hop more nodes learn about the message. Exponential D, disregarding duplicates.
- Difficult to know source of message
- Propagation time depends on:
  - D: connections of each node
  - Number of nodes in the network.
- Problem?
  - Permissionless, so we need rate limit

Background

Zero-knowledge

- Allows to prove a given statement without revealing information.
- Computation intensive: proof generation + verification.

Rate Limiting Nullifiers (RLN)

- Used for rate limiting
- Prove that:
  - Sender knows a secret
  - Hash of secret part of a Merkle tree
  - No double signaling happened

Background

Waku:

- Suite of protocols
- Uses gossipsub for routing
- Uses zero-knowledge proofs (RLN) for rate limiting.
- Principles:
  - Privacy-preserving
  - Decentralized
  - Permissionless
  - Generalized
- But, feasible for real applications?
  - We analyze from latency PoV.

Latency

- **Definition**: Time it takes for a message to propagate to all peers.
- **Goal**: Measure it.

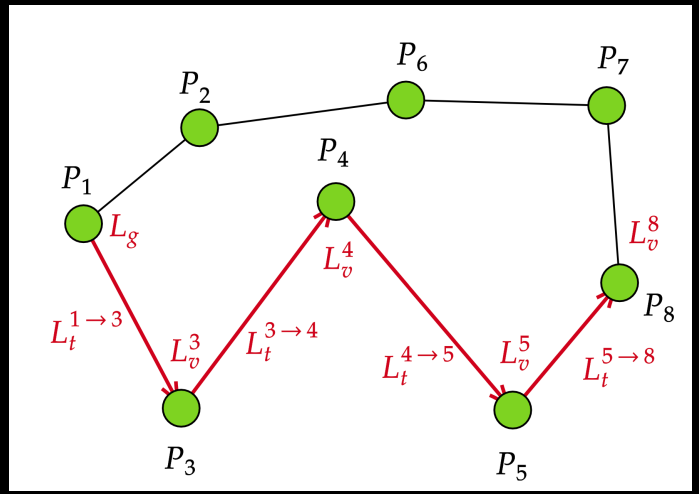

P1 sends 💬 P2. Total latency is:

- L_g: Proof generation
- $L_t^{i\to j}$: Propagation time peer $i \to j$
- L_v^i: Validation time at peer i

Simulations

Single-host simulations

- Mesh of 1000 peers
- D=6: Full mesh connections per peer
- Worst case amount of hops ≈ 4
- Using shadow simulator
  - Simulated latencies (150 ms RTT)
  - Bandwidth constrained (100 Mbps)
  - No CPU time, simulated with await.
- 10 nodes inject a message with a timestamp
- 9900 messages are received at different timestamps (each node receives 10 msg)

Simulations

Single-host simulations

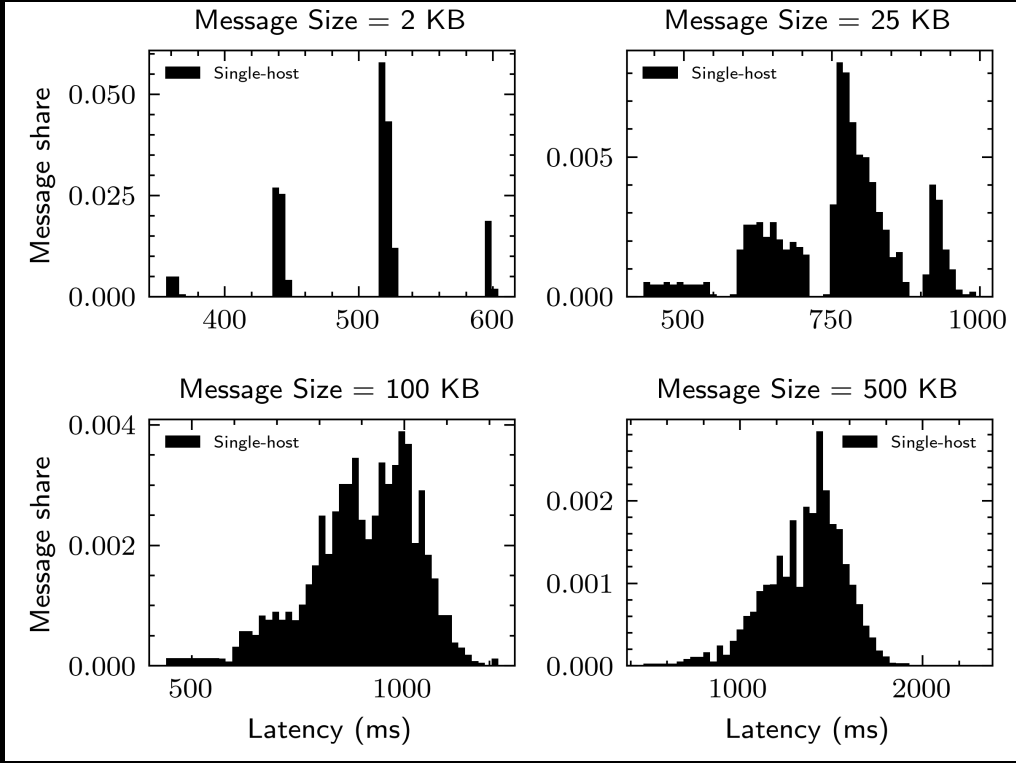

- Small messages (=2KB) latency is RTT/2 scaled by the amount of hops.
- 95% of messages =25KB propagate <1second to all peers
- For big messages (=500KB) < 1.7 seconds with a worst case of 3 seconds.

Simulations

Multi-host simulations

- Real virtual machines (Digital Ocean)
- At different world locations: Singapore, Bangalore, SF, NY and Frankfurt.
- Real Internet conditions.
- Difference from single-host:
  - Just 5 nodes connected in cascade.
  - Cost effective.
  - Boundaries (best/worst case) instead of latency distribution.
  - Same worst amount of hops ≈4

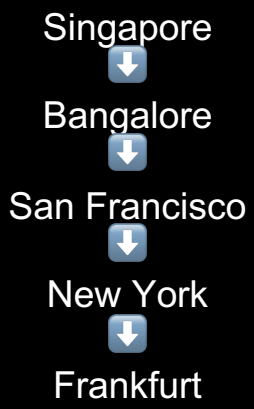

Simulations

Multi-host simulations

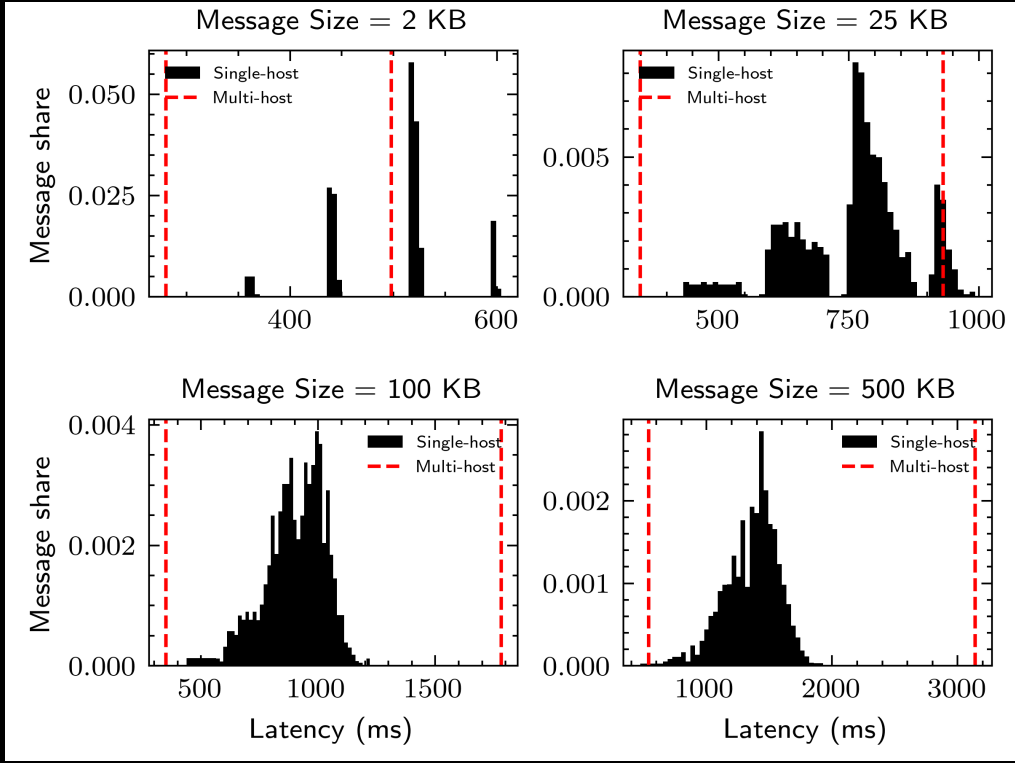

- In black the same distribution as before.
- In red, results from multi-host simulations.
- Validates the single-host simulations (with some nuances)
- See paper for latencies on each hop.
- See paper for RTT between nodes.

Contributions

- Quantify Waku message propagation latencies for different message sizes in a decentralized network (1000 nodes).
  - Best/worst case
  - 95% cases
- Model and identify latency contributions.
- Benchmark RLN implementation and its impact in Waku.
- Validate simulations with multi-host measurements using real network conditions.

Future work

- Dynamic topologies: node churn
- Higher message rates
- Onchain RLN
- Security and privacy analysis
- Compare with other P2P protocols

Questions